

The Core of the Black Canyon Computer Corporation


JOHN COULEUR

The 600 started when I headed the ATLAS digital data recording project in the Heavy Military Equipment Department, Syracuse, NY (HMED). I wanted to design a computer.

I had developed a simple processor to add to the data recording system to do the preflight checkout of the ATLAS missile guidance radar. When GE received a contract to develop the MISTRAM missile tracking radar for testing the Minute Man missile, I upgraded the checkout processor to a more powerful design which could serve as the control center of the MISTRAM tracking system.

The M236, as the new computer was called, was a high speed, 36-bit, minicomputer. It did the down range acquisition, missile tracking, antenna parallax correction, missile position reporting, and system checkout for the MISTRAM system. The Air Force liked the computer, from a standpoint of both cost/performance and reliability, and bought several more for missile tracking applications.

At the same time as I was developing the M236, the computer department, under Claire Lasher was developing a new 32 bit computer family (W, X, Y, & Z). As the design progressed, it became apparent that the upper members of the line would have severe performance, cost, and manufacturing problems. GE management was unhappy with the progress of the project and the high development costs. GE internal large computer users were unhappy with the architecture of the new line and with the prospect of having to replace IBM 7090s with slower computers from the new line.



Barney Oldfield developed a simple processor to add to the data recording system to do the preflight checkout of the ATLAS missile guidance radar. USAF Photo

About a year after the first MISTRAM system was installed in the early 1960s, the Air Force asked me to develop an improved M236 with floating point instructions so they could do down-range impact prediction. In addition, they wanted a Fortran compiler so they could write programs at the site. I asked Don Shell, manager of the mathematics department of the GE Research Laboratory, if he would develop a Fortran compiler for the new version of the M236, which was to be called the M336. Shortly afterward, we wrote a proposal for the Air Force in which the multiprocessor, multimemory, multi-I/O module, architecture was introduced, along with the concept of asynchronous timing connections between the processors, I/O modules, and memories. The architecture was given the M2360 designation. The base/boundary register and slave/master modes of operation were added in late 1963.

After he saw the specifications of the M2360, Shell decided that the M2360 would be a much better large scale computer than the one that Phoenix was developing. As one of the computing leaders in GE, he had influence with the GE large scale computer users, and convinced them that the M2360 computer would make a better upgrade from the IBM 7090 than the computer that Phoenix was developing. GE management did a study of the M2360 and concluded that the new computer would have much better performance, a much lower development cost, and a much lower manufacturing cost than the upper member of the Phoenix line.

The GE 600 Lineage

In early 1963, GE management asked Lasher to consider changing the high end Phoenix computer to the M2360. Claire vetoed the proposed switch and soon was replaced by Harrison Van Aken, who initiated the 600 project; the go-ahead was given on April 1, 1963. The financial plan was such that if every large scale IBM 7090 user within GE replaced their IBM machine with the new GE computer, the 600 project would pay for itself by eliminating the rental costs of the in-house IBM 7090s.

I think that 570 Lexington Avenue wanted Lasher to accept the M2360 and my group on their merits. When he could not, they found someone who could.

Periodically we were called to Phoenix between January and March of 1963 to make proposals and presentations on the M2360. Until the last minute, everyone hoped that the computer department could be persuaded to accept the M2360 and

participate in a joint design effort. The Phoenix technical staff were adamant in their refusal to accept the M2360 or any part of the design or technology. They wanted to continue the Y machine project and apparently felt that the features of the Y would more than make up for the poor performance and high cost, or that they could fix the problems.

Problems of the Y machine were well known within GE. Many GE people associated with computers felt that the Y should be killed and the computer department should be put out of its misery.

Probably, if my group and the M2360 had not come along, GE's patience with the computer department would have run out much earlier and GE's participation in the commercial computer business would have ended in 1963 or 1964. GE was the second biggest IBM mainframe customer behind the U.S. government, and it must have been particularly galling to 570 Lexington Avenue to be paying for both computer department losses and IBM rental at the same time.

I was at Van Aken's company-wide meeting in the Spring of 1963. The GE attendees mostly were IBM customers. They would have been very content to stay with IBM machines (at least at that time). As did all large-scale computer users, they felt that IBM could not be beaten. Of course, they did not consider the fact that GE paid IBM more in rental for their computers than the GE 600 project was estimated to cost. Except for one or two diehards, the tone was "I think that you are out of your mind. Where do I sign up?"

The GE 600 project was set up in Phoenix under Dr. John Weil, a physicist from the Nuclear Department in San Jose, California. I was the head of central system hardware development, and Ed Vance was brought in from Nuclear to head software development. Responsibility for peripherals, input/output, and mechanical design remained with the old Phoenix engineering group. We had one year for development in Syracuse, and I was to move my group, now grown to 20 people, from Syracuse to Phoenix sometime during the next year.

Ed Vance put together a team of GE's large scale computer users to write the specifications for a highly advanced operating system and software family. Most of the team members eventually joined the computer department. Ed filled the rest of the software organization with outside people.

After the project started there was so much fear of contamination of my group by the Phoenix engineers that John Weil kept the project off campus for the first six months, and physically isolated in the plant for a long time after that. There was a considerable fear that the penchant of the Phoenix designers for elegance (translated "putting in everything but the kitchen sink") would creep into the 600 design.

The old Phoenix computer department people still had responsibility for the 225 and 400 lines and peripheral management, as well as responsibility for the design of the 600 input/output subsystem, so very few people were actually put out of a job. The project started out under something less than the best of circumstances.

The GE 600 project had been under way for about a year

when, in the spring of 1964, I invented the translation lookaside buffer (TLB)¹. The TLB was the first practical method of implementing paged/virtual memory. About a month later, the leaders of MIT's Project MAC swung through Phoenix on a search for a platform for the Project MAC timesharing system (Multics).

The GE attendees mostly were IBM customers. They would have been very content to stay with IBM machines (at least at that time). As did all large-scale computer users, they felt that IBM could not be beaten. Of course, they did not consider the fact that GE paid IBM more in rental for their computers than the GE 600 project was estimated to cost.

After the Project MAC people presented their ideas on a computer architecture for supporting timesharing, it became apparent that the use of paging was the only way that their concept could be made to work. I showed them the TLB and how it would interact with their concept of segmentation. MIT liked the idea. GE agreed to modify the 600 design to include paging and segmentation, with MIT (through their ARPA grant) paying for the development. The modified computer was named the GE 645. Bell Laboratories also liked the new system and ordered several 645s.

Defection of two of IBM's biggest and most influential customers stirred the interest of other large computer users who were very unhappy with the 32 bit word length and other features of the IBM System/360. Before the first GE 600 computer was shipped, GE had received letters of intent from more than half of the IBM 7090 computer users in the country.

Except for MIT and Bell Laboratories, all of the sales and shipments were standard GE 600s, though even MIT and Bell Laboratories took early delivery of 600s for their software development. There were, of course, the usual early problems with the hardware. The software was pushing the state of the art and had more severe problems. The first version of GECOS² was too big and too slow, and had to be redesigned. Since early problems were common in the computer industry, the customers were not particularly disturbed. They expected that the problems would be fixed and they would have an early start with what appeared to be the most advanced computer on the market.

GE smelled success, and decided that the computer department was too small an organization to support the anticipated

1. Ted Glaser and John Couleur were granted a patent in 1968 for paging and segmentation memory management; the patents were assigned to MIT and GE respectively. Unknown to Couleur, later MIT granted IBM a royalty-free license to the patent.

2. General Electric Compatible Operating System, c.f. CTSS, Compatible Time-Sharing System, developed by Corbató at MIT.

Black Canyon Computing

level of business. Lou Rader was brought in from UNIVAC to head the newly formed computer group. Lou Wengert was brought in from the Small Motor Division to be the head of the computer division; Ervine M. Koeritz, who was previously the manager of the Precious Metals Department, was brought in as manager of the Computer Equipment Department; Tom Vanderslice was brought in to head the Peripherals Operation. Many others were brought in from GE divisions to fill the many management jobs that were created.

As far as I can remember, very few of the people that were brought in knew much about the computer business except Lou Rader. Most of the new arrivals had no experience in electronics. They were all "GE professional managers."

One customer, who had taken delivery on five of the early systems, summed up the software status with the statement: "GECOS is much better than the IBM System/360 operating system, and far ahead of the System/360 operating system. You have problems, but you know what they are and how to fix them. You are coming out of the woods and IBM is just beginning to find out what their software problems are." Our other customers felt much the same way.

The 645 and Paging³

At the time of the development of the GE 600/645, time-sharing was becoming big. It seemed as if every research group was developing a time-sharing system. MIT had CTSS running on an IBM 7090 and was trying to build a successor. At the time, memory was very expensive (\$500,000 for 256 KB) and time-sharing required large amounts of memory to support many users. Accordingly, it was necessary to swap users in and out of memory to a drum as they completed their time quantum or waited for input from the keyboard.

In the swapping process, the movement of user programs and pieces of user programs (segments) out of memory left irregular sized holes which seldom matched the size of segments to be moved in. The proposed Multics system architecture for Project MAC was particularly sensitive to the problem of moving segments since it encouraged the use of smaller segments and more frequent swapping of segments, and therefore would create a much bigger "hole matching problem."


The use of paging could eliminate the hole matching problem since segments could be broken (mapped) into a number of fixed size pages and all of the holes would be the same size as the pages.

The problem with implementing paging was that each page address had to be relocated (translated) in some manner at a speed which was significantly faster than main memory. If a linearly addressed high speed memory was used for the translation, a translation word had to be provided for every possible processor address page. With the only high speed memory available being registers, this approach was prohibitively expensive for an address space as big as the 600 had.

The lookaside buffer was invented at the University of Manchester and today is known as a cache. At the time it was prohibitively expensive to implement a big enough, fast enough buffer to use as a main memory cache, and therefore the lookaside buffer was never used.

3. The problems described here were initially encountered during the development of the Atlas system at Manchester University.

It worked by having a number of registers which held data from the main memory. Each data register had a connected address register which contained the memory address of the data. A comparator on each of the address registers matched (associated) the contents of the address register with the current memory address. If there was a match, the contents of the attached data register were brought out, much quicker than



The GE 625.

© CBI

they could be retrieved from the corresponding memory location. Of course there was considerable additional logic for housekeeping.

I happened to notice that very few registers (16) would be required to hold the page references for a typical program, and that unlike a linearly addressed

translation memory, only enough associative registers would be required for the translation lookaside buffer (TLB) to accommodate the current memory references. If the program references moved to another part of the program, the TLB could be reloaded. The TLB worked, made paging practical, and allowed Multics to move segments of programs into memory on demand.

Paging with the TLB also allowed the deferment of loading certain pages until they were needed by the program. Although this technique, now called virtual memory, was very inefficient, it allowed execution of very large programs, which otherwise could not be brought into memory. The TLB is still the preferred technique for implementing paging in computers ranging in size from the Intel 486 to mainframes.

Tape Transport Problems

The tape transport used for the GE 600 was a design which had been bought from DECCA in England. The transport turned out to be a disaster. It could not read tapes reliably, and destroyed tapes in the process of reading and writing them. Rather than admit that they had made a mistake buying the DECCA design and arranging to buy tape transports from outside manufacturers, management of the peripheral division insisted that there was nothing wrong with the DECCA mechanism and that it only required a little tweaking. As one customer said when he returned five systems, "I can't run a computer without tapes. I would wait for you to replace the DECCA tapes, but your management tells me that there is nothing wrong with them. As long as they won't admit the tapes don't work, the tapes will never get fixed."

The new disk file that Phoenix was developing for the GE 600 also had problems, but a team from the GE Research Laboratory had isolated the problems and had proposed fixes. Nevertheless, the disk file project was canceled before a suitable replacement was found. GE had no competitive disk file for the next three years.

The Beginning of the End

Starting with the 25th system, the computers could not get through final testing. An investigation showed that Purchasing had bought transistors from another vendor despite the objec-

tion of Engineering. Manufacturing claimed that bid was \$0.10 cheaper than that of Fairchild and had also offered to perform 100% testing on the transistors. Incoming inspection discovered that a large percentage of the Motorola transistors did not meet specifications but nevertheless had been built into the computers.

The superior cost performance of the GE 600 central system was due to a brilliant, unconventional circuit design done by Bob Sullivan, manager of the GE 600 circuit design unit. Circuit delay calculation was based on the statistical performance (i.e. using average rather than worst case performance to determine circuit delay) that could be expected from transistors in a long logic chain. The concept had been proven in the M236 and, in my (and many expert's) mind was correct. The key to achieving the required average performance of the transistors was careful vendor selection and statistical sampling of production batches. If given percentages of the sample did not hit the minimum and maximum speeds (the one-sigma points of the transistor design), as well as meet other criteria, the entire batch was to be rejected.

Sullivan flatly stated that Motorola's process was not capable of producing transistors which met the required performance range. However, by performing 100% checking against the minimum speed, Motorola felt that they could deliver transistors which met the minimum specifications. These transistors would be entirely at the slow end of the range rather than distributed across the performance range as the system design required.

Even with 100% checking to the slowest speed, a large percentage of the Motorola transistors did not meet specifications and should have been returned to Motorola. At the post mortem meeting, it was discovered that a large number of out of specification transistors had been taken out of the return cage and could not be found.

None of the managers sitting on the review board understood statistical sampling, or why it could be more effective than 100% testing. Nor did the fact that the first 25 systems worked with Fairchild transistors, nor that many out of specification transistors must have found their way into manufacturing had any influence. All that management could understand was that we had achieved better cost/performance than the Y machine and that must have been due to cheating on the circuit specifications.

The engineers held over from the old computer department took advantage of the situation to tell anyone who would listen that the 600 circuit design was too dependent on transistor performance, that the system should have been a 32 bit architecture like the IBM 360, and the 600 should be scrapped in favor of some new design that they were more qualified to do. Unfortunately, most of the management had never been associated with electronics, let alone computers, and had no way of evaluating the issues that were being raised.

Bill White was the head of the peripheral equipment division. Bill was appointed to lead the 600 get well program. He determined that the major problem confronting the GE 600 was the circuit design and would probably require a complete redesign of the machine. The tape and disc problems were forgotten for a while. It was over this issue that I left GE and joined University Computing Co. in Dallas. The people criticizing the circuits also reported to John Weil, whose back-

ground was nuclear physics.

Amid all the yelling, Lou Wengert, deputy division general manager for Information Systems Equipment, panicked. He pulled the 600 off the market, threatening to cancel the product. At which point, GE lost their credibility. Most of the cus-

Rather than admit that they had made a mistake buying the DECCA design and arranging to buy tape transports from outside manufacturers, management of the peripheral division insisted that there was nothing wrong with the DECCA mechanism and that it only required a little tweaking. "I can't run a computer without tapes. I would wait for you to replace the DECCA tapes, but your management tells me that there is nothing wrong with them. As long as they won't admit the tapes don't work, the tapes will never get fixed."

tomers defected, having decided that while IBM might not design the best computers, IBM at least knew how to run a computer business.

In all fairness to Lou Wengert, nothing in his career had prepared him for the chaos of a high tech product introduction. Perhaps Lou Rader could have helped, but he was far removed in his headquarters in Virginia. Wengert was getting different advice from various factions, most of whom had their own private agendas. He took what probably appeared to be the best approach: Stop everything until he could figure out what to do.

Nor is it certain that his action in taking the 600 off the market was solely responsible for the loss of customers. Certainly the situation with the tape transports and the disk, and the internal dissension over the 32 bit vs. 36 bit issue cost customer confidence. For whatever the reason, GE never recovered the lost customers or the momentum.

The 32 bit proponents eventually convinced management to create a new, 32 bit product line which was to be developed on a company wide basis. Engineers and planners from Phoenix, Paris, and Milan participated in the specification and design of the new product line.

In the meantime, something else was happening which would have a big impact on the future of the 600. I had left half the M236 design group in Syracuse to form the nucleus of a military computer department. While specifications for the new 32 bit line were being written, the Syracuse group was reimplementing the 600 architecture with newer, faster cir-

Black Canyon Computing

cuits. Shortly after they finished the design, GE moved the Syracuse group to Phoenix.

When GE management saw the anticipated cost of developing the new 32 bit product line, they sold the computer business to Honeywell, who elected to continue designing the line. However, when the Honeywell sales organization took a close look at the combined GE/Honeywell product lines they found that the only product that was available to sell in the near future was the version of the 600 that had been designed in Syracuse. At sales' request, instructions to improve Cobol processing were added to the 600 processor, time-sharing and transaction processing features were added to GECOS, and the price was cut in half. 600 system sales quickly grew from nothing at the time of the merger to 300 to 400 systems per year.

In 1980, when I left Honeywell, the 600 architecture was still the only profitable computer line Honeywell had. The 8000, as the 600 had been renamed (the 600 was enhanced, re-implemented and renamed several times), was making up the losses of the 32 bit product line and providing half of Honeywell's profit. Honeywell was shipping 350 to 400 systems a year, each selling for \$1-4 million, with 90% margin on the central system. The descendants of the 600 probably generated close to \$1 billion a year revenue for Honeywell.

Compagnie des Machines Bull, who bought out Honeywell, is still reimplementing and selling the enhanced 600 architecture. NEC is building super high speed versions of the architecture. GECOS, or GCOS as Honeywell renamed it, has retained much of its original architecture and is still considered a modern operating system.

Multics was very late, very slow, and the system was very expensive. Several customers bought GE 645s for interactive computation, but the 645 and Multics never became the universally used computing system that its creators hoped and expected. As time went on, the 645 architecture became obsolete and the cost of upgrading Multics to a more modern architecture was prohibitive. Eventually, Honeywell abandoned the 645 and Multics. Bell Laboratories became disillusioned early with the Multics project and withdrew. Bell Laboratories developed UNIX. The name was a pun. UNIX avoided many of the technical excesses of Multics and became a moderate success. It is currently being implemented on a descendant of the 600.

Although few people realize it (or are willing to admit it) the Multics architecture lives on in the Intel 386 and its descendants. With the exception of a few additional features that were added to provide compatibility with the 286 and to eliminate indirection that was a part of the 645 architecture, memory management in the 386 is architecturally very close to that of the original GE 645.

Impressions

As for the environment, all I remember is increasing chaos, starting about the time of the reorganization and continuing long after I left in 1967. If I had been asked who was in charge, I would have said "no one." It seemed that major decisions were being made in Phoenix, but there was no top level of management in Phoenix to review them. I think that this absence of high level control probably had more to do with the bad decisions that were made than any individual.

Wengert, Koeritz, Vanderslice, Thompson, Weil, and White, and Wooley from peripherals, made the decisions that I remember. Plus an individual by the name of Dick Benninghausen in Planning who had an extraordinary influence on all of the above people. Possibly they consulted with Lou Rader or Hersher Cross, but I do not remember Cross's or Rader's presence in Phoenix when decisions were being made or debated, or reference to their approval or advice or contribution to the decisions which had such an impact on the business.

All that I remember of the reorganization that created the deputy division managers, was the story that we were getting ready for a big success, the organization was a prototype for a "group," and that as soon as the success arrived we would become a group.

It has been suggested that Lou Wengert went to Phoenix was "to clean up the mess;" this does not match a conversation that I had with him not long after he arrived. He told me that he came to Phoenix expecting to find a smooth running, highly profitable operation. He seemed to be completely shocked by the problems that started showing up not long after he arrived. He made the comment that the problems were nothing like those that he had encountered in Small Motors. Not that Wengert was responsible for the problems, it was just that nobody knew about them at the time of the reorganization.

Professional Management

I guess the secret of making "professional management" work is that the managers plan, organize, integrate, and measure, while the troops who have been with the business and know the business, do the work. As new managers come in, they get to know the troops and who they can trust and not trust. In this way, the business is actually operated by competent people, managers and troops working as a team.

When so many new managers descended on Phoenix at one time, the new organization destroyed the relationships that had made the business successful and which could have worked to find solutions to the problems. No one knew who to trust. The troops had no reputations with the new managers and with that, lost their influence over the business that they had created. The control was taken by the new managers who made decisions based on their newly granted authority. We all lost our influence. That's when the chaos started.

The End

When I saw the new generation of microprocessors, the Motorola 68020 and the Intel 386/486, I concluded that computer design was a dying profession. The computer of the future would sit on a desktop, be based on one of the two chips and the computer business would be dedicated to developing programs for the millions of computers that would be sold. After several years of consulting, I started a company to develop and sell map processing software. I am very happy to be out of the hardware side of the business.

*John Couleur can be contacted at:
4101 N. 65th Street
Scottsdale AZ 85251
Ph: (602) 423-5693
FAX: (602) 970-0386*